

Sparse Hierarchical Solvers with Guaranteed Convergence

Kai Yang*, Hadi Pouransari†, Eric Darve ‡

Abstract

Solving sparse linear systems from discretized PDEs is challenging. Direct solvers have in many cases quadratic complexity (depending on geometry), while iterative solvers require problem dependent preconditioners to be robust and efficient. Approximate factorization preconditioners, such as incomplete LU factorization, provide cheap approximations to the system matrix. However, even a highly accurate preconditioner may have deteriorating performance when the condition number of the system matrix increases. By increasing the accuracy on low-frequency errors, we propose a novel hierarchical solver with improved robustness with respect to the condition number of the linear system. This solver retains the linear computational cost and memory footprint of the original algorithm.

Keywords: sparse, hierarchical, low-rank, elimination, robust preconditioner

1 Introduction

Solving a large-scale sparse linear system

$$Ax = b \tag{1}$$

with $A \in \mathbb{R}^{n \times n}$ is one of the most challenging tasks in scientific computing. A typical feature of the matrix A from discretized partial differential equations (PDEs) is sparsity, which implies that the number of nonzeros per row and per column is small compared with the number of unknowns n . While it is easy to store the linear system in $\mathcal{O}(n)$ memory and multiply A and a vector with $\mathcal{O}(n)$ computational cost, solving $Ax = b$ in $\mathcal{O}(n)$ (or close to $\mathcal{O}(n)$) memory and CPU time remains a challenging issue.

A straightforward Gaussian elimination results in $\mathcal{O}(n^2)$ storage and $\mathcal{O}(n^3)$ computational cost due to the fill-ins introduced. Improved ordering can reduce the fill-ins. For example, nested dissection ordering only takes storage $\mathcal{O}(n \log(n))$ in 2D and $\mathcal{O}(n^{4/3})$ in 3D for simple geometries and discretizations. However, the computational complexity of nested dissection ordering is $\mathcal{O}(n^{3/2})$ in 2D and $\mathcal{O}(n^2)$ in 3D, for most cases.

Instead of directly inverting the sparse matrix, iterative methods seek an approximate solution to (1) within a small number of iterations. Although iterative methods usually require only $\mathcal{O}(n)$ storage, the convergence is highly problem-dependent. Preconditioners are needed in order to achieve fast convergence. There is an extensive literature on developing preconditioners for various applications. Most preconditioners are problem specific and depend on detailed information, such as the geometry and discretization. Although in general more information on the problem results in a more effective solver, in this paper we focus on robust and purely algebraic preconditioners.

Multigrid method [26, 11, 45, 14, 43] is considered as one of the most effective solvers for discretized elliptic PDEs. Based on the discretization of the PDEs, multigrid method efficiently eliminates the error in a multilevel fashion. As an algebraic variant, algebraic multigrid (AMG) [13, 12, 39, 41, 33, 35] does not

*Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA. yangkai.stanford@gmail.com

†Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA. hadip@stanford.edu

‡Institute for Computational and Mathematical Engineering and Department of Mechanical Engineering, Stanford University, Stanford, CA 94305, USA. darve@stanford.edu

require information of the discretization and constructs multilevel structures based on the sparsity pattern of the matrix. AMG have been shown to be efficient for a wide range of applications, especially when the problem is close to elliptic PDEs. In some classical model problems, like the Poisson equation with constant coefficients, AMG achieves linear complexity in memory and time in a way similar to multigrid method. For example, in [33] aggregation-based AMG is shown to be optimal for model Poisson problem under certain assumptions.

Another purely algebraic preconditioner is the incomplete LU (ILU) preconditioner [40]. ILU resembles Gaussian elimination, except that it drops some of the fill-ins to keep the sparsity. ILU has demonstrated its effectiveness in many applications such as convection-diffusion equations. Although it is uncommon that ILU shows optimal convergence, it improves the performance over that of an unpreconditioned iterative solver. Meanwhile, choosing the proper dropping rules for an efficient ILU factorization is a problem-dependent trial and error process.

Low-rank compression is another type of approaches to construct an efficient and accurate preconditioner. This approach was used in fast multipole method (FMM) [23, 19, 34, 46, 20, 38] to accelerate matrix-vector multiplication and then generalized to hierarchical matrices (\mathcal{H} -matrices) [27, 10, 7] for solving dense linear systems from integral equations. There are different types of \mathcal{H} -matrices depending on the low-rank admissibilities and whether the low-rank bases are nested. The hierarchical off-diagonal low-rank (HODLR) matrix [1, 3, 4] assumes the off-diagonal blocks to be low-rank. If the low-rank bases of these off-diagonal blocks can in addition be nested, then this matrix is called hierarchically semi-separable (HSS) [16, 44, 1]. More general than HSS, \mathcal{H}^2 matrices are assumed to have low-rank structure only for well-separated interactions [10], while maintaining nested low-rank bases. In addition to accelerating matrix-vector multiplication, these low-rank techniques have been also applied to approximately factorize the dense matrix and provide direct solvers [7, 2, 18].

Although mostly applied to solving dense linear systems [22, 30, 21, 17, 18], hierarchical matrix techniques have also been applied to solving sparse linear systems, especially those from the discretization of second order elliptic PDEs. For example, the hierarchical LU and Cholesky factorization developed by Hackbusch et al. [25, 27, 10, 6] follow a tree decomposition, recursively factorizing block matrices and forming Schur complements. More recently, hierarchical interpolative factorization [28, 29] was proposed to solve linear systems from the discretization of integral and differential equations based on elliptic operators. Many other approaches based on low-rank properties have also been proposed to solving sparse linear systems [31, 32, 42].

More recently, [37] extended the approaches in [2, 18] to solving sparse linear systems, assuming that the fill-ins during a block Gaussian elimination have low numerical rank for well-separated interactions. An approximate factorization has been introduced via extended sparsification [37, 15], which has time and memory complexity of $\mathcal{O}(n \log^2 1/\epsilon)$ and $\mathcal{O}(n \log 1/\epsilon)$, respectively. Here ϵ is the accuracy of the low-rank approximation. This application results in an accurate preconditioner for the original linear system. Despite the high accuracy of approximate factorizations, the robustness of this type of approaches is subject to the condition number of the original problem. In [37], rather than solving (1) directly, the following equation is solved instead:

$$A_{\mathcal{H}} x_{\mathcal{H}} = b \quad (2)$$

with

$$\|A_{\mathcal{H}} - A\| \leq \epsilon \|A\|. \quad (3)$$

Then the accuracy of the solution is subject to both ϵ and $\kappa(A) = \|A\| \|A^{-1}\|$:

$$\frac{\|x - x_{\mathcal{H}}\|}{\|x\|} \leq C \epsilon \kappa(A),$$

where C is a constant. As the upper bound for the relative accuracy is proportional to the matrix condition number, the accuracy deteriorates when the condition number of matrix grows. This can break the linear complexity as the condition number increases with the number of unknowns in many applications.

In recent work [8, 9], new techniques that alleviate this sensitivity are proposed. The approximate factorization is required to preserve additional vectors, called side constraints, in order to improve the

robustness of the solver. It is proved in [8] that the approximate factorization that additionally preserves low-frequency eigenvectors is spectrally equivalent to the system matrix. [9] also provides adaptive choices of compression parameters for the preconditioner to be spectrally equivalent with constant side constraints. In this paper, we follow similar ideas and propose an improved version of the algorithm in [37]. We note that although the terminology “hierarchical solver” is used both in [37] and [8, 9], these methods are distinct. The distinction was clarified in [18] for example (see Appendix C at the end of this reference). Briefly, [8, 9] is based on a block LU factorization, while [37] is based on a multilevel algorithm similar to AMG.

In this paper, we first formulate the algorithm in [37] in terms of block matrix factorization and analyze the error of the approximate factorization. Then, we present an estimate of the condition number of the preconditioned linear system, which implies that the performance of the hierarchical solver is restricted by its convergence on low-frequency errors. Based on this observation, we enrich the low-rank bases in the matrix compression process to improve the convergence of our hierarchical solver on low-frequency eigenvectors. Numerical examples illustrate that this modification greatly improves the robustness of the solver. In particular, the number of iterations in GMRES becomes nearly independent of the problem size (whereas the condition number increases like $O(1/h^2)$ in the benchmark problems). We use eigenvectors in the numerical examples to validate the theoretical properties. It is computationally expensive however to obtain such eigenvectors. To make this algorithm more practical, the constant vector, which is a computationally cheap alternative to small eigenvectors¹, is instead preserved in hierarchical solvers and investigated in numerical benchmarks in Section 7.

The remainder of this paper is organized as follows. In Section 2 and 3, the factorization phase and solve phase of the hierarchical solver are explained in detail using block matrix notation. In Section 4, we analyze the error of the hierarchical solver due to the low-rank approximations. In Section 5, we analyze the convergence of hierarchical solver when used a preconditioner. Based on these analysis, we propose in Section 6 a new technique to improve the convergence of the hierarchical solver on low-frequency errors. In Section 7, we benchmark the improvement due to the new technique compared with the original hierarchical solver. We verified on these problems that the convergence becomes nearly independent of the problem size. In the appendix we propose some approaches to let the hierarchical solvers approximately preserve vectors of interest.

2 Sparse linear systems

In this section, we introduce basic concepts about sparse linear systems. Throughout the rest of this paper, we assume that A is a symmetric positive definite (SPD) matrix. Then, the condition number of A is defined by the ratio of extreme eigenvalues: $\kappa(A) = \lambda_{\max}(A)/\lambda_{\min}(A)$.

2.1 Graph representation of a sparse matrix

As the matrix $A = (a_{ij})$ is assumed to be symmetric, it corresponds to an undirected graph $G = (V, E)$, where

$$V = \{1, 2, \dots, n\}, \quad E = \{(i, j) | a_{ij} \neq 0\}.$$

A matrix being sparse indicates that it has only a few nonzeros in each row such that the total number of nonzero entries is $\mathcal{O}(n)$. In terms of the graph, sparsity means that each vertex serves as the end point of $\mathcal{O}(n)$ edges. The graph G is also referred to as the *adjacency graph* of the matrix A .

To better leverage the fast block matrix operations of BLAS3 [5], we introduce block partitioning of matrices. Given a partitioning $\{I_i\}_{i=1}^N$ of V s.t.

$$\cup_i I_i = V, \quad I_i \cap I_j = \emptyset, \text{ if } i \neq j,$$

¹For brevity, we will use the expression “small” eigenvector, to mean the eigenvector associated with the small eigenvalues.

let $A_{I_i I_j}$ denote a block matrix of A with row indices in I_i and column indices in I_j . A new graph with less vertices, denoted by G_b , is then defined as $G_b = (V_b, E_b)$ with

$$V_b = \{1, 2, \dots, N\}, \quad E_b = \{(i, j) | A_{I_i I_j} \text{ has nonzero entries}\}.$$

3 Hierarchical solver

Although the Gaussian elimination works for any linear system, it destroys the sparsity when applied to sparse linear systems by introducing additional nonzero entries, or *fill-ins*, into the factorized matrices L and U . Many attempts have been made to reduce this fill-in via suitable approximation; for example, incomplete LU factorization (ILU) strictly controls the number of fill-ins introduced. Our approach is similar to the ILU in the sense that we also improve the algorithm by reducing the flops and storage spent on the fill-ins. The major difference between our method and ILU is that instead of simply dropping the fill-ins as is the case for ILU, our algorithm compresses the fill-ins based on the low-rank property, simultaneously preserving efficiency and accuracy. A key observation is that the fill-ins corresponding to well-separated interactions are low-rank [37]. For brevity, we simply use *h-solver* to denote this type of hierarchical solvers. In the rest of this section, we introduce the algorithm in [37] using block matrix notation. This is necessary for introducing the improved version of the algorithm in the next few sections.

Given a sparse linear system, the first step of the proposed algorithm is to partition the unknowns into local clusters. The purpose of clustering is to accelerate the algorithm using BLAS3 operations [5]. For example, the partitioning can be formed as is shown in Fig. 1, where each color corresponds to one cluster. This block partitioning corresponds to a new graph with fewer vertices. An example of the new graph is given in Fig. 1. As we will pairwise agglomerate the clusters, we assume that there are $2N$ such clusters, denoted by $\{I_i^r\}_{i=1}^{2N}$. Each cluster I_i^r is called *red node* in [37] and we continue to use this name to distinguish them from other types of clusters.

Before we start the factorization, we first agglomerate the red nodes pairwise. The motivation of this step is to facilitate the low-rank compression and will be explained in more detail later in this paper. Let $\{I_i^s\}_{i=1}^N$ be the clusters formed by

$$I_i^s = I_{2i-1}^r \cup I_{2i}^r, \quad i = 1, 2, \dots, N.$$

The new clusters $\{I_i^s\}_{i=1}^N$ are called *super nodes* in [37].




Figure 1: Partitioning of the graph of a matrix resulting from a 2D discretization. Left: partitioning of the graph into local clusters denoted by different colors. Right: The corresponding graph for the block version of the underlying matrix.

3.1 Adjacency and well-separated interactions

Note that adjacency is purely based on the original graph before any elimination. During the elimination, new edges are introduced, but the adjacency between the super nodes remains unchanged throughout the elimination process.

The interactions between clusters are the off-diagonal blocks in the matrix. For example, the interaction between cluster I_i and I_j ($i \neq j$) is $A_{I_i I_j}$. This rule applies to both red nodes and super nodes and defines the adjacency.

Although the adjacency stays the same, the interactions between clusters change during elimination. Before we eliminate the first super node, the interactions are only between adjacent/neighbor clusters. After we eliminate some clusters, interactions between unadjacent clusters might be introduced into the matrix as fill-ins.

The key idea of [37], as well as those of many other related papers (cite more), is closely related to the concept of **well-separated interactions**, which is borrowed from the context of the fast multipole method (FMM). In FMM, two clusters are **well-separated** if their distance is larger than a fraction of their diameters. A key finding in the FMM is that the interactions between well-separated clusters can be approximated accurately by low-rank matrices, assuming that their interactions are represented by non-oscillatory (“smooth”) kernels [23, 19, 46]. As the h-solver is purely algebraic, it does not have access to the geometric information as the FMM does. Instead, well-separated interactions have to be defined purely based on the graph of the system matrix. A straightforward approach is to define that two clusters are well-separated if they are not neighbors. Although different from the criterion in FMM, this type of well-separated interactions is expected to exhibit similar low-rank properties as those emerging from integral equations. The low-rank properties have been numerically verified in [37].

3.2 Factorization

The factorization phase of the h-solver is similar to that of a Gaussian (or in this instance Cholesky) elimination. The major difference between an h-solver and a Gaussian elimination is that an h-solver first compresses well-separated interactions before it eliminates the current node.

For each super node, first compress the well-separated interactions. Then introduce additional variables and obtain an extended system matrix. After the extension, the current super node only has interactions with neighboring nodes; namely, there is no well-separated interaction between the current super node and the rest of the nodes. Then perform the elimination. We will elaborate on the compression, extension, and elimination steps of the h-solver. Assume that a given ordering of the super nodes is provided: $\{s^1, s^2, \dots, s^N\}$.

Let’s consider the case when we are eliminating a super node s^i . We partition the unknowns of the linear system into three groups, corresponding to

- the current super node s^i , denoted by subscript “s”,
- the neighboring nodes of s^i , denoted by “n”,
- the well-separated nodes from s^i , denoted by “w”.

Then the linear system (1) has the following block form

$$\begin{pmatrix} A_{ss} & A_{sn} & A_{sw} \\ A_{ns} & A_{nn} & A_{nw} \\ A_{ws} & A_{wn} & A_{ww} \end{pmatrix} \begin{pmatrix} x_s \\ x_n \\ x_w \end{pmatrix} = \begin{pmatrix} b_s \\ b_n \\ b_w \end{pmatrix}. \quad (4)$$

In case s^i is not the first super node in the list, the super nodes that are already eliminated are not included in (4). Moreover, new nodes will be introduced into the system in the extension step. Therefore, among the neighboring and well-separated clusters of s^i , there are newly introduced clusters. We do not elaborate on these new clusters here, but just assume the adjacency between s^i and these new clusters are well defined. We will define this adjacency in the extension step.

Note that when dealing with the first super node, there is no nonzero well-separated interactions; namely, $A_{sw} = 0$. A_{sw} becomes nonzero during the algorithm due to fill-ins.

Compression

The well-separated interactions A_{sw} and A_{ws} can be approximated well by low-rank matrices. For example, we perform a truncated Singular Value Decomposition (SVD)

$$A_{sw} = U\Sigma V^T + E_{sw}. \quad (5)$$

Here U and V are both orthonormal, Σ is a diagonal matrix consisting of the most significant singular values of A_{sw} . The number of columns of U is much smaller than the number of rows or columns of A_{sw} . Moreover, E_{sw} is small and bounded by the given compression threshold. For convenience, we define $R^T = \Sigma V^T$. Thus,

$$A_{sw} = UR^T + E_{sw}. \quad (6)$$

We now solve the following approximate linear system

$$\begin{pmatrix} A_{ss} & A_{sn} & UR^T \\ A_{ns} & A_{nn} & A_{nw} \\ RU^T & A_{wn} & A_{ww} \end{pmatrix} \begin{pmatrix} x_s \\ x_n \\ x_w \end{pmatrix} = \begin{pmatrix} b_s \\ b_n \\ b_w \end{pmatrix}. \quad (7)$$

Extension

In order to avoid more and more fill-ins, two additional unknowns, y_b and y_r , are introduced. Then we can instead solve

$$\begin{pmatrix} A_{ss} & A_{sn} & & U & & \\ A_{ns} & A_{nn} & A_{nw} & & & \\ & A_{wn} & A_{ww} & & R & \\ U^T & & & R^T & -I & \end{pmatrix} \begin{pmatrix} x_s \\ x_n \\ x_w \\ y_b \\ y_r \end{pmatrix} = \begin{pmatrix} b_s \\ b_n \\ b_w \\ 0 \\ 0 \end{pmatrix}. \quad (8)$$

as x_s, x_n, x_w in (8) solves (7). In [37], the unknowns y_b and y_r are considered as a black node and a red node. We will explain the underlying meanings of these new nodes at the end of this section. Note that the newly introduced red node y_r will not be merged or eliminated until all of the super nodes $\{s^1, s^2, \dots, s^N\}$ are eliminated.

The red node r needs well-defined neighbor list in order to continue the algorithm. In fact, r inherits the neighbors of s , which may contain super nodes and red nodes.

Elimination

In the elimination step, we eliminate x_s and y_b . First, the linear system is rearranged as follows

$$\begin{pmatrix} A_{ss} & U & A_{sn} & & & \\ U^T & & & & -I & \\ A_{ns} & & A_{nn} & A_{nw} & & \\ & & A_{wn} & A_{ww} & R & \\ & -I & & R^T & & \end{pmatrix} \begin{pmatrix} x_s \\ y_b \\ x_n \\ x_w \\ y_r \end{pmatrix} = \begin{pmatrix} b_s \\ 0 \\ b_n \\ b_w \\ 0 \end{pmatrix}. \quad (9)$$

Denote (9) by

$$K\tilde{x} = \tilde{b}.$$

We eliminate the current super node x_s and the black node y_b : $K = LK_2L^T$ with

$$L = \begin{pmatrix} I & & & & \\ L_{bs} & I & & & \\ L_{ns} & L_{nb} & I & & \\ & & & I & \\ & & S^{-1} & & I \end{pmatrix}$$

$$K_2 = \begin{pmatrix} A_{ss} & & & & \\ & -S & & & \\ & & \tilde{A}_{nn} & A_{nw} & L_{nb} \\ & & A_{wn} & A_{ww} & R \\ & & L_{bn} & R^T & S^{-1} \end{pmatrix} = \begin{pmatrix} A_{ss} & & & & \\ & -S & & & \\ & & & & A^{(i+1)} \end{pmatrix}.$$

Here

$$\begin{aligned} L_{bs} &= U^T A_{ss}^{-1} & L_{ns} &= A_{ns} A_{ss}^{-1} \\ L_{nb} &= A_{ns} A_{ss}^{-1} U S^{-1} & S &= U^T A_{ss}^{-1} U \\ \tilde{A}_{nn} &= A_{nn} - A_{ns} A_{ss}^{-1} A_{sn} + A_{nb} S^{-1} A_{bn}. \end{aligned}$$

In order to solve

$$K_2 L^T \tilde{x} = L^{-1} \tilde{b}, \quad (10)$$

we need to solve a smaller linear system $A^{(i+1)}x = y$. By picking the next super node s^{i+1} , this linear system is again in the form of (4). We can repeat the compression, extension, and elimination procedure.

Note that each time these three steps are repeated, a super node s^i is eliminated and a red node, denoted by r^i , is added to the system. The black node is eliminated right after being introduced to the system. By induction, we can conclude that at the beginning of each round of compression-extension-elimination, the nodes in the system are

$$\{s^i, s^{i+1}, \dots, s^N, r^1, \dots, r^{i-1}\}.$$

By the end of each round, the nodes in the system are

$$\{s^{i+1}, \dots, s^N, r^1, \dots, r^i\}.$$

This process is repeated until all the super nodes on the current level have been eliminated, by which time the nodes in the system are $\{r^1, r^2, \dots, r^N\}$. We denote the resulting matrix by A^{parent} . Then we can again pairwise merge the red nodes in order to form super nodes, and repeat the compression-extension-elimination for all the new super nodes until there is only one red node in the system. Then we factorize the matrix exactly. This procedure is summarized in a recursive fashion by Algorithm 1. We use A^{leaf} to denote the matrix in (1).

Note that as we repeatedly agglomerate pairs of red nodes into super nodes, we assume that the block partitioning of the original linear system (1) forms $N = 2^L$ red nodes, where L is an integer. The assumption can be easily satisfied, for example, by recursively bisecting of the graph of (1).

In [37], this factorization algorithm is explained using hierarchical trees. We briefly explain the connection here and refer interested readers to [37] for more details.

The red nodes are pairwise merged to form super nodes (Fig. 2), and each super node generate a black node and a red node. This procedure reduce the number of red nodes by half. We put each black node on top of the corresponding super node, and then put the new red node on the corresponding black node. After eliminating all the super nodes and the black nodes, we only have the newly introduced red nodes, which are only half as many as the red nodes we started with. As this procedure will be repeated until there is only one red node, all of the nodes actually form a hierarchical tree. (See Fig. 3.)

The tree notation of [37] provides a vivid interpretation of the purely algebraic algorithm of the h-solver. We refer to [37] for more details on the tree formulations.

Algorithm 1 Setup phase of the h-solver

```

SETUP( $A^{\text{leaf}}$ )
function SETUP( $A$ )
  for each pair of sibling red nodes  $r_1$  and  $r_2$  in the red node list do
    Merge  $r_1$  and  $r_2$  into a super node  $s$ .
  end for
  for  $s$  in the super node list do
    Compress the well-separated interaction  $A_{sn}, A_{ns}$ . See (7).
    Introduce a black node  $b$  and a parent red node  $r$ . Extend the matrix. See (8).
    Eliminate the super node  $s$  and the black node  $b$ . See (10).
  end for
  if The parent level is the root then
    Factorize  $A^{\text{parent}}$  directly.
  else
    Setup( $A^{\text{parent}}$ ).
  end if
end function
  
```




Figure 2: Merge sibling red nodes




Figure 3: An example of hierarchical trees

Computational complexity of the factorization

It is shown in [37] that the factorization algorithm has complexity $\mathcal{O}(nd_{l_{\max}}^2)$, as long as there is a constant α such that $d_i/d_{i+1} < \alpha < 2^{1/3}$, $\forall 0 \leq i \leq l_{\max} - 1$. Here n is the number of unknowns and d_i is the maximum number of unknowns in a super node (or red node) on the i -th level. The l_{\max} -th level is the leaf level. The ratio d_i/d_{i+1} depends on the low-rank property of the well-separated interactions. For example, given the red node size d_{i+1} on level $i + 1$, the super node has size $2d_{i+1}$. Suppose that the truncated SVD is able to provide the required accuracy with rank less than d_{i+1} , then the red node on level i has size $d_i \leq d_{i+1}$. Then, d_i/d_{i+1} satisfies the criterion. We refer to [37] for the details of complexity analysis.

3.3 Solve

After factorizing the matrix, we approximately solve $Ax = b$ with a given right hand side b . This procedure can be denoted as solving the linear system $A_{\mathcal{H}}x_{\mathcal{H}} = b$.

Based on the factorization, the solution process is also recursive. Starting with a right hand side b and let $A^{(0)} = A$. If $i < N$, we eliminate a super node from $A^{(i)}$. For this, we split the right hand side b as

$$b = \begin{pmatrix} b_s \\ b_n \\ b_w \end{pmatrix}.$$

We have to solve the extended system (10) in three steps.

1. Forward substitution: $\phi = L^{-1}\tilde{b}$. Component-wise we have

$$\phi = \begin{pmatrix} \phi_s \\ \phi_b \\ \phi_n \\ \phi_w \\ \phi_r \end{pmatrix} = \begin{pmatrix} b_s \\ -U^T A_{ss}^{-1} b_s \\ b_n - A_{ns} A_{ss}^{-1} b_s + A_{ns} A_{ss}^{-1} U S^{-1} U^T A_{ss}^{-1} b_s \\ b_w \\ S^{-1} U^T A_{ss}^{-1} b_s \end{pmatrix}.$$

2. Recursive solve: $\psi = K_2^{-1}\phi$. This step involves inverting $A^{(i+1)}$. Component-wise we have:

$$\psi = \begin{pmatrix} \psi_s \\ \psi_b \\ \psi_n \\ \psi_w \\ \psi_r \end{pmatrix} = \begin{pmatrix} A_{ss}^{-1} b_s \\ S^{-1} U^T A_{ss}^{-1} b_s \\ \psi_n \\ \psi_w \\ \psi_r \end{pmatrix} \quad \text{and} \quad A^{(i+1)} \begin{pmatrix} \psi_n \\ \psi_w \\ \psi_r \end{pmatrix} = \begin{pmatrix} \phi_n \\ \phi_w \\ \phi_r \end{pmatrix}.$$

As for solving the second equation above, we go to the next super node and apply the algorithm in case $i + 1 < N$. If $i + 1 = N$, we go to the next group of super nodes, which are on the next upper level in the hierarchical tree.

3. Backward substitution: $\tilde{x} = L^{-T}\psi$. Component-wise we have

$$\tilde{x} = \begin{pmatrix} x_s \\ y_b \\ x_n \\ x_w \\ y_r \end{pmatrix} = \begin{pmatrix} A_{ss}^{-1}(b_s - U S^{-1}(U^T A_{ss}^{-1}(b_s - A_{sn}\psi_n) - \psi_r) - A_{sn}\psi_n) \\ S^{-1}(U^T A_{ss}^{-1}(b_s - A_{sn}\psi_n) - \psi_r) \\ \psi_n \\ \psi_w \\ \psi_r \end{pmatrix}.$$

Then the solution is

$$x = \begin{pmatrix} x_s \\ x_n \\ x_w \end{pmatrix}.$$

Here we introduce some more notation to simplify the expressions:

$$\begin{aligned} P_{ss} &:= S^{-1}U^T A_{ss}^{-1} \\ M_{ss} &:= A_{ss}^{-1} - P_{ss}^T S P_{ss}. \end{aligned}$$

Then some of the expressions can be simplified:

$$\begin{aligned} b_n - A_{ns}A_{ss}^{-1}b_s + A_{ns}A_{ss}^{-1}US^{-1}U^T A_{ss}^{-1}b_s &= b_n - A_{ns}M_{ss}b_s \\ A_{ss}^{-1}(b_s - US^{-1}(U^T A_{ss}^{-1}(b_s - A_{sn}\psi_n) - \psi_r) - A_{sn}\psi_n) &= M_{ss}(b_s - A_{sn}\psi_n) + P_{ss}^T\psi_r. \end{aligned}$$

Note that P_{ss} maps the unknowns of the current super node to the unknowns of the corresponding parent red node. M_{ss} is closely related to A_{ss}^{-1} . In fact, when U vanishes, M_{ss} is exactly A_{ss}^{-1} . On the other hand, when U is invertible, $M_{ss} = 0$. M_{ss} can be factored as:

$$M_{ss} = (I - A_{ss}^{-1}US^{-1}U^T)A_{ss}^{-1} = P_{\text{perp}}A_{ss}^{-1},$$

where P_{perp} is an A_{ss} orthogonal projection onto $\text{span}\{U\}^\perp$. Here $\text{span}\{U\}$ denote the space spanned by the columns of U .

The solve algorithm involves recursion from one super node to the next super node. However, it is possible to rearrange the order of computation to first finish the computation for all of the super nodes on one level of the hierarchical tree and then recurse to the next upper level of the tree. The solve step is formulated as Algorithm 2. It makes a recursive call when moving on to the next super node. If $i = N$, then all of the current super nodes are eliminated and the system matrix $A^{(N)}$ corresponds to parent level red nodes. Then, call the solve algorithm recursively on the parent level.

We name this h-solver **IFMM** (*Inverse Fast Multipole Method*), in order to distinguish it from the improved h-solver to be presented later in this paper.

4 Error Estimate

In this section, we consider how the local truncated SVDs contribute to the error of $A_{\mathcal{H}}$. Note that in this paper we discuss the truncated SVDs, but other methods such as rank-revealing QR, rank-revealing LU, ACA could be used.

We will only consider a two-level method; namely, on the parent level the linear system is solved exactly. There are N steps of compression and elimination, where N is the number of super nodes on the leaf level.

We first recognize that the extension step could be expressed as a matrix product.

$$\begin{aligned} &\begin{pmatrix} A_{ss} & A_{sn} & UR^T & & \\ A_{ns} & A_{nn} & A_{nw} & & \\ RU^T & A_{wn} & A_{ww} & & \\ & & & -I & \\ & & & & -I \end{pmatrix} \\ &= \begin{pmatrix} I & & & U \\ & I & & \\ & & I & R \\ & & & I \\ & & & & I \end{pmatrix} \begin{pmatrix} A_{ss} & A_{sn} & & U \\ A_{ns} & A_{nn} & A_{nw} & \\ & A_{wn} & A_{ww} & R \\ U^T & & & -I \\ & R^T & -I & \end{pmatrix} \begin{pmatrix} I & & & \\ & I & & \\ & & I & \\ & & R^T & I \\ U^T & & & I \end{pmatrix} \end{aligned} \quad (11)$$

In order to avoid constantly changing the size of the linear system due to the extension steps, we consider the extended system which contains all of the super nodes, the black nodes and the parent red nodes. The

Algorithm 2 Solve phase of h-solver

$x = \text{SOLVE}(A^{\text{leaf}}, b^{\text{leaf}})$

function SOLVE(A, b)

For all parent level node p , set $r_p = 0$.

for s in the super node list **do**

Forward substitution for the neighbors n of s :

$$\phi_n = r_n - A_{ns} M_{ss} r_s$$

end for

for r in the parent level red node list **do**

Restriction:

$$\phi_r = \phi_r + P_{ss} \phi_s$$

end for

Parent level solve:

$$\psi^{\text{parent}} = \text{SOLVE}(A^{\text{parent}}, \phi^{\text{parent}})$$

for r in the parent level red node list **do**

Prolongation:

$$x_s = P_{ss}^T \psi_r$$

end for

for s in the super node list of a reversed order **do**

Backward substitution for the neighbors n of s :

$$x_s = x_s + M_{ss}(b_s - A_{sn} x_n)$$

end for

return x

end function

following equation is equivalent to (1) in the sense that x_s in its solution solves (1). Note that the solid lines in this equation separate the three different types of nodes.

$$\left(\begin{array}{c|c|c} A & & \\ \hline & & -I \\ \hline & -I & \end{array} \right) \begin{pmatrix} x_s \\ x_b \\ x_r \end{pmatrix} = \begin{pmatrix} b \\ 0 \\ 0 \end{pmatrix}. \quad (12)$$

Denote the matrix in (12) by \mathcal{K}_1 . Then for each super node, the compression, extension and elimination steps can be reformulated in the extended system. Consider the i -th super node. Right before the compression, the matrix is denoted by \mathcal{K}_i .

In the compression step, we decompose the well-separated interactions into the low-rank part and error part, as is shown in (6). We embed the error part E_{sw} and E_{ws} in the extended system with all types of nodes and denote it by \mathcal{E}_i . Let \mathcal{K}_i^- denote the matrix \mathcal{K}_i with the well-separated interactions A_{sw} and A_{ws} replaced by low-rank matrices UR^T and RU^T , respectively. Thus, $\mathcal{K}_i = \mathcal{E}_i + \mathcal{K}_i^-$.

For the extension step, we first introduce the matrix

$$\mathcal{U}_i = \left(\begin{array}{c|c|c} I & R_i & U_i \\ \hline & I & \\ \hline & & I \end{array} \right),$$

where the only nonzero parts of R_i and U_i are R and U , respectively. For example, R is located on the rows corresponding to the nodes well-separated from the i -th super node, and on the columns corresponding to the i -th black node. Similarly, we can locate U in U_i . Then the extension step is denoted by $\mathcal{K}_i^- = \mathcal{U}_i \mathcal{K}_i^E \mathcal{U}_i^T$. Here \mathcal{K}_i^E corresponds to the extended matrix in the second line of (11).

For the elimination step, we first embed L in the extended system, with 1's on the rest of the diagonal entries, and denote it by \mathcal{L}_i . Then the elimination is denoted by $\mathcal{K}_i^E = \mathcal{L}_i \mathcal{K}_{i+1} \mathcal{L}_i^T$.

Therefore, the compression, extension, and elimination steps for the i -th super node is denoted by

$$\mathcal{K}_i = \mathcal{E}_i + \mathcal{U}_i \mathcal{L}_i \mathcal{K}_{i+1} \mathcal{L}_i^T \mathcal{U}_i^T.$$

If we recursively apply this equation, the factors $\mathcal{U}_i \mathcal{L}_i$ and $\mathcal{L}_i^T \mathcal{U}_i^T$ are applied on the error \mathcal{E}_j for any $i < j \leq N$. Note that $\mathcal{L}_i \mathcal{E}_{i+1} \mathcal{L}_i^T = \mathcal{E}_{i+1}$, because \mathcal{E}_{i+1} is zero in the rows and columns corresponding to the first i super nodes and black nodes. Multiplying \mathcal{U}_i with \mathcal{E}_{i+1} may create additional nonzeros in the matrix. The reason is that \mathcal{E}_{i+1} may include nonzero interaction between the $(i+1)$ -th super node and the i -th parent red node, which we denote by E_{ps} . In fact, the additional error term introduced by this multiplication is $U E_{ps}$. As U is the orthonormal low-rank basis, the additional error satisfies $\|U E_{ps}\|_F = \|E_{ps}\|_F$.

Moreover, note that $\mathcal{L}_{i-1} \mathcal{U}_i \mathcal{E}_{i+1} \mathcal{U}_i^T \mathcal{L}_{i-1}^T = \mathcal{U}_i \mathcal{E}_{i+1} \mathcal{U}_i^T$. Then we can obtain the following equation

$$\mathcal{K}_1 = \sum_{i=1}^N \tilde{\mathcal{E}}_i + \mathcal{U}_1 \mathcal{L}_1 \cdots \mathcal{U}_N \mathcal{L}_N \mathcal{K}_{N+1} \mathcal{L}_N^T \mathcal{U}_N^T \cdots \mathcal{L}_1^T \mathcal{U}_1^T$$

with

$$\tilde{\mathcal{E}}_i = \begin{cases} \mathcal{U}_1 \cdots \mathcal{U}_{i-1} \mathcal{E}_i \mathcal{U}_{i-1}^T \cdots \mathcal{U}_1^T, & i \geq 2, \\ \mathcal{E}_1, & i = 1. \end{cases}$$

As each \mathcal{U}_i introduces an additional error for the i -th super node, the total error can be estimated

$$\|\tilde{\mathcal{E}}_i\|_F^2 \leq 2\|\mathcal{E}_i\|_F^2.$$

Let $\mathcal{K} = \mathcal{K}_1$, $\mathcal{E} = \sum_i \tilde{\mathcal{E}}_i$, and $\mathcal{K}_{\mathcal{H}} = \mathcal{U}_1 \mathcal{L}_1 \cdots \mathcal{U}_N \mathcal{L}_N \mathcal{K}_{N+1} \mathcal{L}_N^T \mathcal{U}_N^T \cdots \mathcal{L}_1^T \mathcal{U}_1^T$. $\mathcal{K}_{\mathcal{H}}$ is extended hierarchical factorization, while \mathcal{K} is basically the original linear system, with auxiliary variables introduced. Assume that we apply an absolute error tolerance $\|\mathcal{E}_i\|_F \leq \epsilon$. Then it is straightforward to obtain the following estimate

$$\|\mathcal{K} - \mathcal{K}_{\mathcal{H}}\|_F \leq (2N)^{1/2} \epsilon.$$

Although the hierarchical solver is realized by solving the extended linear system

$$\mathcal{K}_{\mathcal{H}} \begin{pmatrix} x_s \\ x_b \\ x_r \end{pmatrix} = \begin{pmatrix} b \\ 0 \\ 0 \end{pmatrix}, \quad (13)$$

it could also be formulated as a linear system with the same number of unknowns as (2). After all, the hierarchical solver is a linear mapping from a right hand side b to a solution $x_{\mathcal{H}}$.

Consider a general extended linear system:

$$\begin{pmatrix} A_- & B^T \\ B & C \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} b \\ 0 \end{pmatrix}. \quad (14)$$

If the x in the solution of the extended system (14) solves the original system (1) for any b , we say that (14) is an **equivalent extension** of (1). The following lemma provides a necessary and sufficient condition for an equivalent extension.

Lemma 1. *Assume that A and C are invertible. Then (14) is an equivalent extension of (1) if and only if*

$$A_- - B^T C^{-1} B = A$$

holds.

Introduce the block partitioning, where “f” denotes the super nodes and “c” denote the black nodes and the parent red nodes. Thus, we have the following partitioning:

$$\mathcal{K} = \begin{pmatrix} K_{ff} & K_{fc} \\ K_{cf} & K_{cc} \end{pmatrix} \quad \text{and} \quad \mathcal{E} = \begin{pmatrix} E_{ff} & E_{fc} \\ E_{cf} & \end{pmatrix}.$$

As (12) is an equivalent extension of (1), and (13) is an equivalent extension of (2), the following equations hold based on Lemma 3

$$\begin{aligned} A &= K_{ff} + E_{ff} - (K_{fc} + E_{fc})K_{cc}^{-1}(K_{cf} + E_{cf}) \\ A_{\mathcal{H}} &= K_{ff} - K_{fc}K_{cc}^{-1}K_{cf}. \end{aligned}$$

Therefore, we can quantify the error

$$A - A_{\mathcal{H}} = E_{ff} - E_{fc}K_{cc}^{-1}K_{cf} - K_{fc}K_{cc}^{-1}E_{cf} - E_{fc}K_{cc}^{-1}E_{cf},$$

and obtain the following error estimate.

Theorem 1. *The following inequality holds*

$$\|A - A_{\mathcal{H}}\|_F \leq \sqrt{1 + 2\|K_{cc}^{-1}K_{cf}\|_F^2} \|\mathcal{E}\|_F + \|K_{cc}^{-1}\|_F \|\mathcal{E}\|_F^2.$$

Moreover, assume that we have N leaf super nodes and that each compression based on (5) satisfies that $\|E_{sw}\| \leq \epsilon$. Then

$$\|A - A_{\mathcal{H}}\|_F \leq (2N)^{1/2} \sqrt{1 + 2\|K_{cc}^{-1}K_{cf}\|_F^2} \epsilon + 2N\|K_{cc}^{-1}\|_F \epsilon^2.$$

5 Convergence Analysis

In this section, we analyze the convergence of the hierarchical solvers we have introduced. Approximation properties in the form of (3) are assumed and the conclusion applies to a wide range of approximate factorization solvers. In fact, [8] has a similar convergence analysis. They assume that the approximate operator preserves vectors that are close to the small eigenvectors. Then they prove that the condition number of the preconditioned system $\kappa(A_{\mathcal{H}}^{-1}A)$ is independent of the problem size. In this section, we present a similar but more general convergence analysis.

Assume that a given collection of subsets of \mathbb{R}^n : $\{S_i\}_{i=1}^N$ satisfies the following conditions.

(A1) $\{S_i\}$ form a direct sum of \mathbb{R}^n ; namely, $\mathbb{R}^n = S_1 \oplus S_2 \oplus \cdots \oplus S_N$. In addition, $\{S_i\}$ are A -orthogonal; namely, give any $i \neq j$, $v_i^T A v_j = 0$ holds for $\forall v_i \in S_i, v_j \in S_j$.

(A2) $A_{\mathcal{H}}$ approximates A up to relative error ϵ_i on S_i , i.e.

$$\|(A - A_{\mathcal{H}})v\| \leq \epsilon_i \|A\| \|v\|, \quad \forall v \in S_i.$$

Note that **(A1)** is possible when S_i 's are spanned by nonoverlapping subsets of eigenvectors of A . As A is SPD, we use λ_{max} , the largest eigenvalue of A , instead of $\|A\|$ in the upper bound of **(A2)**.

Note that we can get an upper bound of $\kappa(A_{\mathcal{H}}^{-1}A)$ by finding constants c_0 and c_1 such that

$$c_0(Av, v) \leq (A_{\mathcal{H}}v, v) \leq c_1(Av, v), \quad \forall v \in \mathbb{R}^n.$$

We have the following estimates for c_1 and c_2 :

$$c_1 = \max_v \frac{(A_{\mathcal{H}}v, v)}{(Av, v)} = 1 + \max_v \frac{((A_{\mathcal{H}} - A)v, v)}{(Av, v)} \quad (15)$$

$$c_0 = \min_v \frac{(A_{\mathcal{H}}v, v)}{(Av, v)} = 1 - \max_v \frac{((A - A_{\mathcal{H}})v, v)}{(Av, v)}. \quad (16)$$

Define $P_i = S_i \oplus \cdots \oplus S_N$. Then for $v_i \in P_i$, we can decompose it as $v_i = v_{i+1} + \tilde{v}_i$, with $v_{i+1} \in P_{i+1}$ and $\tilde{v}_i \in S_i$.

$$\begin{aligned} & \max_{v_i \in P_i} \frac{((A_{\mathcal{H}} - A)v_i, v_i)}{(Av_i, v_i)} \\ &= \max_{\tilde{v}_i \in S_i, v_{i+1} \in P_{i+1}} \frac{((A_{\mathcal{H}} - A)(v_{i+1} + \tilde{v}_i), v_{i+1} + \tilde{v}_i)}{(A(v_{i+1} + \tilde{v}_i), v_{i+1} + \tilde{v}_i)} \\ &\leq \max_{\tilde{v}_i \in S_i, v_{i+1} \in P_{i+1}} \frac{((A_{\mathcal{H}} - A)v_{i+1}, v_{i+1})}{(Av_{i+1}, v_{i+1}) + (A\tilde{v}_i, \tilde{v}_i)} + \max_{\tilde{v}_i \in S_i, v_{i+1} \in P_{i+1}} \frac{((A_{\mathcal{H}} - A)(2v_{i+1} + \tilde{v}_i), \tilde{v}_i)}{(A(v_{i+1} + \tilde{v}_i), v_{i+1} + \tilde{v}_i)} \\ &\leq \max_{v_{i+1} \in P_{i+1}} \frac{((A_{\mathcal{H}} - A)v_{i+1}, v_{i+1})}{(Av_{i+1}, v_{i+1})} + \max_{\tilde{v}_i \in S_i, v_{i+1} \in P_{i+1}} \frac{\|(A_{\mathcal{H}} - A)\tilde{v}_i\| \|2v_{i+1} + \tilde{v}_i\|}{(A(v_{i+1} + \tilde{v}_i), v_{i+1} + \tilde{v}_i)} \\ &\leq \max_{v_{i+1} \in P_{i+1}} \frac{((A_{\mathcal{H}} - A)v_{i+1}, v_{i+1})}{(Av_{i+1}, v_{i+1})} + \max_{\tilde{v}_i \in S_i, v_{i+1} \in P_{i+1}} \frac{\epsilon_i \lambda_{max} \|\tilde{v}_i\| \|2v_{i+1} + \tilde{v}_i\|}{(A(v_{i+1} + \tilde{v}_i), v_{i+1} + \tilde{v}_i)}. \end{aligned}$$

Define

$$\mu_i := \min_{v_i \in P_i} \frac{(Av_i, v_i)}{(v_i, v_i)}, \quad 1 \leq i \leq N.$$

μ_i measures the lowest frequency of vectors in P_i . For example, if P_i is spanned by the eigenvectors e_i, e_{i+1}, \dots, e_N , then $\mu_i = \lambda_i$.

In order to estimate

$$\max_{\tilde{v}_i \in S_i, v_{i+1} \in P_{i+1}} \frac{\|\tilde{v}_i\| \|2v_{i+1} + \tilde{v}_i\|}{(A(v_{i+1} + \tilde{v}_i), v_{i+1} + \tilde{v}_i)},$$

we simply assume that neither $\|\tilde{v}_i\|$ nor $\|2v_{i+1} + \tilde{v}_i\|$ is zero. Then we have the following estimate:

$$\begin{aligned}
& \frac{(A(v_{i+1} + \tilde{v}_i), v_{i+1} + \tilde{v}_i)}{\|\tilde{v}_i\| \|2v_{i+1} + \tilde{v}_i\|} \\
&= \frac{(A(2v_{i+1} + \tilde{v}_i), 2v_{i+1} + \tilde{v}_i) + 3(A\tilde{v}_i, \tilde{v}_i)}{4\|\tilde{v}_i\| \|2v_{i+1} + \tilde{v}_i\|} \\
&= \frac{\|2v_{i+1} + \tilde{v}_i\|}{4\|\tilde{v}_i\|} \frac{(A(2v_{i+1} + \tilde{v}_i), 2v_{i+1} + \tilde{v}_i)}{\|2v_{i+1} + \tilde{v}_i\|^2} + \frac{\|\tilde{v}_i\|}{4\|2v_{i+1} + \tilde{v}_i\|} \frac{3(A\tilde{v}_i, \tilde{v}_i)}{\|\tilde{v}_i\|^2} \\
&\geq \mu_i \left(\frac{\|2v_{i+1} + \tilde{v}_i\|}{4\|\tilde{v}_i\|} + \frac{3\|\tilde{v}_i\|}{4\|2v_{i+1} + \tilde{v}_i\|} \right) \\
&\geq \frac{\sqrt{3}}{2} \mu_i.
\end{aligned}$$

Therefore, the following inequality holds

$$\max_{v_i \in P_i} \frac{((A_{\mathcal{H}} - A)v_i, v_i)}{(Av_i, v_i)} \leq \max_{v_{i+1} \in P_{i+1}} \frac{((A_{\mathcal{H}} - A)v_{i+1}, v_{i+1})}{(Av_{i+1}, v_{i+1})} + \frac{2\epsilon_i \lambda_{max}}{\sqrt{3}\mu_i}.$$

We can repeat this procedure for $1 \leq i \leq N-1$ and get

$$\max_{v_i \in P_i} \frac{((A_{\mathcal{H}} - A)v_i, v_i)}{(Av_i, v_i)} \leq \frac{\epsilon_N \lambda_{max}}{\mu_N} + \sum_{i < N} \frac{2\epsilon_i \lambda_{max}}{\sqrt{3}\mu_i}.$$

Note that for $A - A_{\mathcal{H}}$, the same bound can be derived; namely,

$$\max_{v_i \in P_i} \frac{((A - A_{\mathcal{H}})v_i, v_i)}{(Av_i, v_i)} \leq \frac{\epsilon_N \lambda_{max}}{\mu_N} + \sum_{i < N} \frac{2\epsilon_i \lambda_{max}}{\sqrt{3}\mu_i}.$$

Based on (15), we get

$$c_1 = 1 + \frac{\epsilon_N \lambda_{max}}{\mu_N} + \frac{2}{\sqrt{3}} \sum_{i < N} \frac{\epsilon_i \lambda_{max}}{\mu_i}.$$

In addition, assume ϵ_i 's are small enough such that

$$\frac{\epsilon_N \lambda_{max}}{\mu_N} + \frac{2}{\sqrt{3}} \sum_{i < N} \frac{\epsilon_i \lambda_{max}}{\mu_i} < 1. \tag{17}$$

Then we get

$$c_0 = 1 - \frac{\epsilon_N \lambda_{max}}{\mu_N} - \frac{2}{\sqrt{3}} \sum_{i < N} \frac{\epsilon_i \lambda_{max}}{\mu_i}.$$

We now have an upper bound for the condition number $\kappa(A_{\mathcal{H}}^{-1}A)$, as detailed in the theorem below.

Theorem 2. Assume (A1), (A2) and (17). Then, the following inequality holds

$$\kappa(A_{\mathcal{H}}^{-1}A) \leq \frac{1 + \epsilon_N \lambda_{max}/\mu_N + \frac{2}{\sqrt{3}} \sum_{i < N} \epsilon_i \lambda_{max}/\mu_i}{1 - \epsilon_N \lambda_{max}/\mu_N - \frac{2}{\sqrt{3}} \sum_{i < N} \epsilon_i \lambda_{max}/\mu_i}.$$

Example 1. Assume we have a uniform accuracy ϵ for all eigenvectors. Then it is the special case $\mathbb{R}^n = S_1$ and, therefore, $\mu_1 = \lambda_{min}$, the minimum eigenvalue of A . Then we have

$$\kappa(A_{\mathcal{H}}^{-1}A) \leq \frac{1 + \epsilon \lambda_{max}/\lambda_{min}}{1 - \epsilon \lambda_{max}/\lambda_{min}} = \frac{1 + \epsilon \kappa(A)}{1 - \epsilon \kappa(A)}.$$

This upper bound indicates that, with a uniform accuracy ϵ for error components of all frequency, ϵ has to decrease as fast as $1/\kappa(A)$ in order to keep $\kappa(A_{\mathcal{H}}^{-1}A)$ bounded independent of $\kappa(A)$.

Example 2. Let v_1, v_2, \dots, v_m be a few vectors and $S_1 := \text{span}\{v_i\}_{i=1}^m$. Let $S_2 := S_1^{\perp A}$ be the A -orthogonal complement of S_1 in \mathbb{R}^n . Then,

$$\kappa(A_{\mathcal{H}}^{-1}A) \leq \frac{1 + \frac{2}{\sqrt{3}}\epsilon_1\kappa(A) + \epsilon_2\lambda_{\max}/\mu_2}{1 - \frac{2}{\sqrt{3}}\epsilon_1\kappa(A) - \epsilon_2\lambda_{\max}/\mu_2}$$

with

$$\mu_2 = \min_{v \in S_1^{\perp A}} \frac{(Av, v)}{(v, v)}. \quad (18)$$

In particular, we can pick $0 < \eta < 1$ in order to get a partitioning of small and large eigenvalues:

$$S_1 = \text{span}\{e_i\}_{i:\lambda_i < \eta}, \quad S_2 = \text{span}\{e_i\}_{i:\lambda_i \geq \eta}.$$

Then $\mu_2 \geq \eta$. Therefore,

$$\kappa(A_{\mathcal{H}}^{-1}A) \leq \frac{1 + \frac{2}{\sqrt{3}}\epsilon_1\kappa(A) + \epsilon_2\lambda_{\max}/\eta}{1 - \frac{2}{\sqrt{3}}\epsilon_1\kappa(A) - \epsilon_2\lambda_{\max}/\eta}.$$

This example implies a practical strategy of improving the original hierarchical solver. We consider S_1 to be spanned by a few low-frequency eigenvectors and improve the accuracy on these eigenvectors. For the rest of the eigenvectors in S_2 we maintain the accuracy of the original hierarchical solver algorithm. Note that in general with η being constant, the dimension of S_1 increases with the dimension of A .

Example 3. Consider eigenvectors e_1, e_2, \dots, e_n and the corresponding eigenvalues $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. Assume that $S_i := \text{span}\{e_i\}$, $1 \leq i \leq n$. Then $P_i = \text{span}_{j \geq i}\{e_j\}$ and $\mu_i = \lambda_i$. Therefore, the following inequality holds

$$\kappa(A_{\mathcal{H}}^{-1}A) \leq \frac{1 + \frac{2}{\sqrt{3}} \sum_{i < n} \epsilon_i \lambda_{\max}/\lambda_i + \epsilon_n \lambda_{\max}/\lambda_n}{1 - \frac{2}{\sqrt{3}} \sum_{i < n} \epsilon_i \lambda_{\max}/\lambda_i - \epsilon_n \lambda_{\max}/\lambda_n}.$$

This upper bound suggests that the accuracy on each eigenvector shall be proportional to the magnitude of the corresponding eigenvalue in order to be efficient.

6 Improved hierarchical solver with guaranteed convergence

Based on the convergence analysis, we can see that, in order to guarantee the convergence, we need to improve the accuracy of the h-solver on low-frequency error components. There are two approaches for this purpose. The first approach is to let the $A_{\mathcal{H}}$ preserves the action of A on a few low-frequency vectors, such that the h-solver is exact on these error components. The second approach is to improve the accuracy of $A_{\mathcal{H}}$ on a few low-frequency vectors. Both of these approaches add additional bases in the truncated SVD. In this section, we introduce the first approach. The second approach is covered in Appendix A.

Note that in [8] modifications to the hierarchical Cholesky decomposition are presented for the solver to satisfy the side constraint. Although we have the same goal in this paper, the algorithmic details are different due to the distinction in the hierarchical factorization algorithms.

Based on Example 2, being exact on a few eigenvectors, i.e. letting $\epsilon_1 = 0$, leads to the following estimate

$$\kappa(A_{\mathcal{H}}^{-1}A) \leq \frac{1 + 2\epsilon_2\lambda_{\max}/\eta}{1 - 2\epsilon_2\lambda_{\max}/\eta}.$$

In this case if we can in addition guarantee constant η and the same accuracy ϵ_2 for problems of different sizes, then $\kappa(A_{\mathcal{H}}^{-1}A)$ will be uniformly bounded independent of problem sizes. However, in general this would be difficult to achieve. For example, for the 2D Poisson equation discretized by the finite difference method on uniform grid, the number of eigenvalues smaller than a constant $\eta > 0$ grows linearly with the number of unknowns. Nevertheless, it is still promising that preserving a few eigenvectors may improve the performance of the hierarchical solvers.

Next, we consider how a vector is preserved during compression, extension, and elimination.

Preserving a vector during compression

During the compression step, only two off-diagonal block matrices are modified. It is straightforward to see that the modified blocks need to preserve the action of the original blocks on the corresponding segments of the vector. The following algorithm provides sufficient conditions for preserving vectors during the compression.

Lemma 2. *Given*

$$\mathcal{A} = \begin{pmatrix} S & B^T & C^T \\ B & D & E^T \\ C & E & F \end{pmatrix} \quad \text{and} \quad \phi = \begin{pmatrix} \phi_x \\ \phi_y \\ \phi_z \end{pmatrix},$$

let $\tilde{\mathcal{A}}$ be an approximation of \mathcal{A} by replacing B and B^T by \tilde{B} and \tilde{B}^T , respectively. Assume that $B^T \phi_y = \tilde{B}^T \phi_y$ and $B \phi_x = \tilde{B} \phi_x$ hold. Then $\mathcal{A} \phi = \tilde{\mathcal{A}} \phi$.

Based on Lemma 2, the compression (6) has to satisfy

$$A_{sw} \phi_y = U R^T \phi_y \quad \text{and} \quad A_{ws} \phi_x = R U^T \phi_x. \quad (19)$$

In order to satisfy this property, the compression of A_{sw} has to follow a more complicated approach, instead of a straightforward SVD. In general, we are looking for a decomposition of the form

$$A_{sw} = U R^T + \hat{U} \hat{R}^T = [U, \hat{U}] \begin{bmatrix} R^T \\ \hat{R}^T \end{bmatrix},$$

where the concatenated matrix $[U, \hat{U}]$ is orthonormal. Then a sufficient condition for (19) to hold is

$$\phi_x \in \text{span}\{U\} \quad \text{and} \quad A_{sw} \phi_y \in \text{span}\{U\}. \quad (20)$$

Therefore, we first orthonormalize the following matrix using QR or SVD

$$[\phi_x, A_{sw} \phi_y] = U_1 R_1^T,$$

where the columns of U_1 are orthonormal. We keep the component of A_{sw} that is in the range of U_1 , i.e. $U_1 K_1^T$, with $K_1^T := U_1^T A_{sw}$. Then, perform a low-rank factorization of the following matrix

$$(I - U_1 U_1^T) A_{sw} = U_2 K_2^T + E.$$

Here $U_2 K_2^T$ has low rank but approximates $(I - U_1 U_1^T) A_{sw}$ well in the sense that $\|E\|/\|A_{sw}\|$ is less than the given compression parameter ϵ . There are various approaches to obtain this approximation, such as the truncated SVD, the rank-revealing QR and LU, ACA, etc. Then we obtain the following approximate factorization:

$$A_{sw} \approx [U_1, U_2] \begin{bmatrix} K_1^T \\ K_2^T \end{bmatrix}.$$

We can check that

$$\|A_{sw} - [U_1, U_2] \begin{bmatrix} K_1^T \\ K_2^T \end{bmatrix}\| \leq \epsilon \|A_{sw}\|.$$

Preserving a vector in the extension step

During the extension step, the size of the linear system changes. We need to consider carefully which vectors to preserve when the number of unknowns changes.

Given the current global matrix \mathcal{A} and the vector ϕ we want to preserve, consider the following extended linear system

$$\begin{pmatrix} \mathcal{A}_- & \mathcal{K}^T \\ \mathcal{K} & \mathcal{M} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} b \\ 0 \end{pmatrix},$$

the solution of which is consistent with that of $\mathcal{A}x = b$ under the assumption $\mathcal{A} = \mathcal{A}_- - \mathcal{K}^T \mathcal{M}^{-1} \mathcal{K}$. For brevity we denote the extended system by $\mathcal{A}_E x_E = b_E$. When \mathcal{A}_E is further approximated by

$$\tilde{\mathcal{A}}_E := \begin{pmatrix} \tilde{\mathcal{A}}_- & \tilde{\mathcal{K}}^T \\ \tilde{\mathcal{K}} & \tilde{\mathcal{M}} \end{pmatrix}$$

in future steps, we also get an approximation for \mathcal{A} : $\tilde{\mathcal{A}} := \tilde{\mathcal{A}}_- - \tilde{\mathcal{K}}^T \tilde{\mathcal{M}}^{-1} \tilde{\mathcal{K}}$. We need to determine the vector to be preserved by the extended system, such that $\tilde{\mathcal{A}}\phi = \mathcal{A}\phi$.

The following lemma provides us a guideline.

Lemma 3. *Given ϕ , let $\xi = -\mathcal{M}^{-1} \mathcal{K}\phi$ and $\phi_E = \begin{pmatrix} \phi \\ \xi \end{pmatrix}$. Then a sufficient condition for $\tilde{\mathcal{A}}\phi = \mathcal{A}\phi$ to hold is $\tilde{\mathcal{A}}_E \phi_E = \mathcal{A}_E \phi_E$.*

Proof. Given

$$\begin{pmatrix} \tilde{\mathcal{A}}_- & \tilde{\mathcal{K}}^T \\ \tilde{\mathcal{K}} & \tilde{\mathcal{M}} \end{pmatrix} \begin{pmatrix} \phi \\ \xi \end{pmatrix} = \begin{pmatrix} \mathcal{A}_- & \mathcal{K}^T \\ \mathcal{K} & \mathcal{M} \end{pmatrix} \begin{pmatrix} \phi \\ \xi \end{pmatrix} = \begin{pmatrix} \mathcal{A}\phi \\ 0 \end{pmatrix},$$

we solve for ξ using the second row: $\xi = -\tilde{\mathcal{M}}^{-1} \tilde{\mathcal{K}}\phi$. Then, we have

$$\mathcal{A}\phi = \tilde{\mathcal{A}}_- \phi + \tilde{\mathcal{K}}^T \xi = \tilde{\mathcal{A}}_- \phi - \tilde{\mathcal{K}}^T \tilde{\mathcal{M}}^{-1} \tilde{\mathcal{K}}\phi = \tilde{\mathcal{A}}\phi.$$

□

Although these equations seem complicated, it turns out that in the context of our algorithm, many of these terms simplify. In fact, the only calculation that needs to be done is $U^T \phi_x$, which is the representation (in the basis U) of the vector ϕ we wish to preserve for the current cluster x . $U^T \phi_x$ is then the vector to be preserved at the parent level (the red nodes right above the leaf level in the tree). Let's explain this step-by-step.

Based on Lemma 3, in order to preserve

$$\phi = \begin{pmatrix} \phi_s \\ \phi_n \\ \phi_w \end{pmatrix},$$

we need to first calculate

$$\xi = - \begin{pmatrix} 0 & -I \\ -I & 0 \end{pmatrix}^{-1} \begin{pmatrix} U^T & 0 & 0 \\ 0 & 0 & R^T \end{pmatrix} \begin{pmatrix} \phi_s \\ \phi_n \\ \phi_w \end{pmatrix} = \begin{pmatrix} R^T \phi_w \\ U^T \phi_s \end{pmatrix}.$$

The term $R^T \phi_w$ is associated with the black node, while $U^T \phi_s$ is for the parent red node. For the extended system \mathcal{A}_E , the extended vector to be preserved is then

$$\phi_E = \begin{pmatrix} \phi_s \\ R^T \phi_w \\ \phi_n \\ \phi_w \\ U^T \phi_s \end{pmatrix}.$$

Now let's eliminate the leaf node and its black node. These are two steps in the block Cholesky factorization.

Preserving a vector during elimination

Recall that the elimination process is

$$K = LK_2L^T.$$

Given an approximation \tilde{K}_2 of K_2 , the matrix \tilde{K}

$$\tilde{K} := L\tilde{K}_2L^T$$

is our approximation of K .

Lemma 4. *Given a vector ϕ , a sufficient condition for $K\phi = \tilde{K}\phi$ to hold is*

$$K_2\psi = \tilde{K}_2\psi, \quad \text{where } \psi = L^T\phi.$$

The vector ψ is of the form

$$\psi = L^T\phi = \begin{pmatrix} \tilde{x}_s \\ \tilde{y}_b \\ x_n \\ x_w \\ y_r \end{pmatrix}, \quad \text{given that } \phi = \begin{pmatrix} x_s \\ y_b \\ x_n \\ x_w \\ y_r \end{pmatrix}.$$

Here $\tilde{\cdot}$ denotes an update resulting from the multiplication by L^T . However, at that point in the algorithm the only blocks that need compressing are those involving x_n , x_w and y_r . The unknowns for s and b have already been eliminated so that \tilde{x}_s and \tilde{y}_b are automatically preserved in the rest of the factorization.

In summary, looking back at all the steps, we simply proceed as follows. When we need to compress well-separated blocks, we calculate the vectors to preserve: ϕ_x and $A_{sw}\phi_y$ (see (20)) and then form the appropriate U and R matrices. We also calculate $U^T\phi_x$, which becomes the vector associated with the parent red node and needs to be preserved during subsequent block low-rank approximations.

During the factorization each block matrix preserves the corresponding local segment of the global vectors to be preserved. A consequence is that the approximate operator $A_{\mathcal{H}}$ preserves a space of vectors with dimension much larger than the number of vectors to be preserved. For example, assume that we want $A_{\mathcal{H}}$ to preserve k linearly independent vectors and there are N super nodes on the leaf level. Then, $A_{\mathcal{H}}$ preserves a vector space of dimension Nk . This fact was observed and explained as well in [8]. This property will help $A_{\mathcal{H}}$ to better approximate the low-frequency eigenvectors.

It is also possible to preserve the vectors up to a given accuracy. Details on this approach are provided in Appendix A.

We name this improved version of h-solver **GC**, which is short for *Guaranteed Convergence*.

Computational complexity of the GC

Although in the GC we incorporate additional bases in the compression step, the criterion for linear complexity is still $d_i/d_{i+1} \leq \alpha < 2^{1/3}$, $\forall 0 \leq i \leq l_{\max} - 1$. Note that d_i consists of both the additional vectors incorporated and the bases from the truncated SVD.

7 Numerical tests

In this section, we present numerical tests to demonstrate the performance of the h-solver and validate the theories proposed in previous sections. The implementation is based on the LoRaSp² package. LoRaSp

²**Low Rank Sparse** solver, bitbucket.org/hadip/lorasp.

depends on Eigen [24] for basic linear algebra calculations and SCOTCH [36] for graph partitioning. In the numerical tests, we solve the d-dimensional Poisson equation (d=2, 3)

$$\begin{cases} -\nabla \cdot (\alpha \nabla u) = f & \text{in } \Omega := [0, 1]^d \\ u = 0 & \text{on } \partial\Omega \end{cases}$$

discretized by the second-order central finite difference method. We solve both 2D and 3D Poisson equations. As we use standard five-point (2D) or seven point (3D) stencil, we can make use of the geometry information and define two clusters as well-separated clusters if their distance is larger than the size of the clusters.

In the following numerical tests, we first compare the performance of the improved h-solver (GC) with that of the original h-solver (IFMM) for a constant-coefficient Poisson equation, i.e., $\alpha(x) = 1$ in Ω . Note that as we are using the standard five-point (2D) or seven-point (3D) stencil, the eigenvalues and eigenvectors of the system matrices are available. Therefore, we are able to enrich the GC with the smallest eigenvector. In practice, computing eigenvectors might be computationally expensive. Therefore, we also enrich the hierarchical basis with the constant vector, which approximates the low-frequency vectors well for the Poisson equation.

In the second part of the numerical tests, we also solve a variable-coefficient Poisson equation. As the eigenvectors do not have explicit form for these cases, we only compare the IFMM with the GC-constant.

7.1 Constant-coefficient Poisson equation

2D cases

In the 2D test, we first compare h-solvers by using them in the stationary iteration. Given an initial guess x_0 , the iterative scheme $x_{k+1} = x_k + A_{\mathcal{H}}^{-1}(b - Ax_k)$ produces a sequence x_0, x_1, \dots . As we randomly set up the solution x^* and then obtain the right hand side $b = Ax^*$, the error of the k -th iteration can be evaluated as $e_k = x_k - x^*$. The iteration converges when the relative error $\|e_k\|/\|e_0\|$ is less than 10^{-6} .

In Table 1, we tested a sequence of problems with increasing number of unknowns n . The tree depths are adjusted accordingly to let the size of each leaf red node be 8. Table 1 shows the number of the stationary iteration for the relative error to decrease to less than 10^{-6} . In addition to the IFMM, we also tested the new h-solver enriched with a constant vector (GC-constant) and with the smallest eigenvector (GC-eigenvector).

Table 1: Number of the stationary iteration of the h-solvers for the constant coefficient Poisson equation. Leaf red nodes have a size of 8. The stopping criterion is relative error less than 10^{-6} . The compression parameter ϵ is 0.1 for all cases.

n	tree depth	IFMM	GC-constant	GC-eigenvector
$(2^5)^2$	7	4	3	3
$(2^6)^2$	9	5	4	4
$(2^7)^2$	11	12	5	4
$(2^8)^2$	13	24	4	5
$(2^9)^2$	15	79	5	5
$(2^{10})^2$	17	289	5	6

From Table 1, the number of iterations of the IFMM increases as the number of unknowns grows. In comparison, the GC-constant and the GC-eigenvector have an almost constant number of iterations for all cases.

In practice, preconditioned Krylov subspace iterations are much more robust than the stationary iteration. In Table 2, we show the number of GMRES iterations preconditioned by the h-solvers. For each of the h-solvers, compression parameters $\epsilon = 0.1, 0.2$, and 0.3 are tested.

	tree depth	IFMM			GC-constant			GC-eigenvector		
$n \setminus \epsilon$		0.1	0.2	0.3	0.1	0.2	0.3	0.1	0.2	0.3
$(2^5)^2$	7	5	6	7	5	6	7	5	6	7
$(2^6)^2$	9	6	8	10	6	7	8	6	7	8
$(2^7)^2$	11	8	11	14	7	8	10	6	8	10
$(2^8)^2$	13	11	17	22	7	9	11	7	9	11
$(2^9)^2$	15	18	30	39	7	10	14	8	10	14
$(2^{10})^2$	17	31	53	71	8	11	16	8	11	16

Table 2: Number of GMRES preconditioned by h-solvers solving constant coefficient Poisson equation. The tolerance ϵ used in the h-solver is shown on the second row. Leaf red nodes have a size of 8. The stopping criterion is relative residual less than 10^{-10} .

For all of the h-solvers, increasing ϵ results in a larger number of iterations. For each value of ϵ , the IFMM shows a significantly increasing number of iterations. The number of iterations of the GC-constant and the GC-eigenvector, however, are almost constant and only increase mildly with the number of unknowns. We used only relatively large values for ϵ since this leads already to a small number of iterations. ϵ below 0.1 reduces the number of iterations further. However, it leads to a more expensive preconditioner so that overall there is no gain.

In both the stationary iteration and preconditioned GMRES, the performance of the GC-constant and the GC-eigenvector are almost identical. This indicates that the constant vector is a good alternative for the smallest eigenvector for the 2D Poisson equation.

3D cases

For the 3D Poisson equation, the condition number of the system matrices grows less rapidly with respect to the number of unknowns than in the 2D cases, although the storage cost and CPU time per iteration become higher due to denser system matrices.

In Table 3, we show the number of GMRES iterations when preconditioned by the IFMM and the GC with $\epsilon = 0.2$ and 0.3 . Moreover, as in [9], adaptive choice of ϵ for different problem sizes has been proposed. We also apply this adaptive strategy in order to explore its influence on the h-solvers. In [9], four approaches are tested:

- $\mathcal{H}\text{Chol}$: constant compression parameter for all cases. It is comparable to the IFMM with constant ϵ .
- $\mathcal{H}\text{Chols}$: adaptive compression parameter $\epsilon_l = \epsilon h d_l$, where d_l is the minimum diameter of the finite element basis functions on level l . It is comparable to the IFMM with adaptive ϵ . As we use finite difference and regular partitioning, we have $d_l = 2^{(l_{\max}-l)/3}h$. If we use this criterion, ϵ_l becomes too small and is clearly sub-optimal for our method. Therefore, we adopt a different compression parameter $\epsilon_l = \epsilon h 2^{(l_{\max}-l)/3}$ (we drop a factor h) such that the computation for large problems is more affordable. With this criterion, the relative accuracy between leaves and root is the same as in [9] (large ϵ at the root), but the scaling with h is $O(h)$ instead of $O(h^2)$.
- $\mathcal{M}\mathcal{H}\text{Chol}$: preserving the constant vector and constant compression parameter. This is comparable to the GC-constant.
- $\mathcal{M}\mathcal{H}\text{Chols}$: preserving the constant vector and adaptive compression parameter $\epsilon_l = \epsilon h D_l^{-1}$ (D_l : maximum diameter at level l). This is comparable to the GC with adaptive $\epsilon_l = \epsilon 2^{(l-l_{\max})/3}$ (small ϵ at the root). The compression parameter for the leaf level remains constant when the number of unknowns changes.

tree depth		IFMM				GC-constant			
$n \setminus \epsilon$		0.2	0.3	$\epsilon_{l_{\max}} = 8/\sqrt[3]{n}$		0.2	0.3	$\epsilon_{l_{\max}} = 0.2$	$\epsilon_{l_{\max}} = 0.3$
8^3	6	5	5	7		5	5	5	5
16^3	9	6	6	8		5	6	5	6
32^3	12	7	8	9		6	7	6	7
64^3	15	8	11	9		6	8	6	8

Table 3: Number of GMRES iterations with different hierarchical preconditioners solving constant coefficient Poisson equation. Leaf red nodes have a size of 8. The stopping criterion is a relative residual less than 10^{-10} . In the IFMM, $\epsilon_{l_{\max}}$ means that we use a small ϵ_l at the leaf and a constant ϵ_l at the root (see $\mathcal{H}\text{Chols}$). In the GC, ϵ_l is constant at the leaf (as indicated) and goes down to $O(h)$ at the root (see $\mathcal{MH}\text{Chols}$).

From Table 3, the IFMM has a number of iterations growing very slowly with the number of unknowns, even though large values of ϵ (0.2, 0.3) are used. This is quite different from [9] where they report a much faster increase in the number of iterations with problem size. The adaptive choice of ϵ leads to a roughly constant number of iterations. We note that the adaptive $\epsilon_{l_{\max}}$ keeps decreasing as n increases, so that the cost of the preconditioner goes up with n at a faster rate than n . If we assume that the rank is $O(\log 1/\epsilon)$, then the CPU time grows roughly like $O(n \log^2 n)$. A similar increase in CPU time occurs in [9].

The GC-constant improves over the IFMM. In this case, the adaptive strategy leads to larger ranks near the root. However, since the number of iterations is already nearly constant, this choice for ϵ_l does not seem to make a difference. The CPU time is larger with adaptive ϵ_l compared to non-adaptive. So the adaptive choice does not seem preferable for this benchmark.

Different results might be obtained on larger grids. However, limitations in memory do not allow us to run larger test cases. The largest case we ran is $n \approx 262k$.

7.2 Variable-coefficient Poisson equation

Variable-coefficient Poisson equation is frequently encountered in applications of engineering problems, for example, the heat conduction problem with more than one medium. Another example is fluid flow in heterogeneous porous media. We consider the following two variable-coefficient cases.




Figure 4: Subdomains for piecewise-constant coefficient Poisson equation

- **Piecewise-constant coefficient:** Given $\Omega = [0, 1] \times [0, 1]$, define a subdomain $\Omega_1 := [1/4, 3/4] \times [1/4, 3/4]$. The rest of the domain is $\Omega_2 := \Omega \setminus \Omega_1$. (See Fig. 4.) Let

$$\alpha(x) = \begin{cases} 10^{-5}, & \text{if } x \in \Omega_1, \\ 1, & \text{if } x \in \Omega_2. \end{cases}$$

- **Random coefficient:** we use a simplified model (not physical) in which we use uniformly distributed random variables in $[0, 1]$ for α in the discretized system.

tree depth		IFMM			GC-constant		
$n \setminus \epsilon$		0.1	0.2	0.3	0.1	0.2	0.3
$(2^5)^2$	7	6	6	6	7	8	7
$(2^6)^2$	9	6	7	8	7	8	9
$(2^7)^2$	11	9	10	14	8	10	12
$(2^8)^2$	13	11	16	21	10	10	13
$(2^9)^2$	15	17	29	37	9	11	15
$(2^{10})^2$	17	31	54	70	10	13	18

Table 4: Number of GMRES iterations preconditioned by h-solvers for the piecewise-constant coefficient Poisson equation. Leaf red nodes have a size of 8. The stopping criterion is relative residual less than 10^{-10} .

Table 5: Number of GMRES iterations preconditioned by h-solvers for the random coefficient Poisson equation. Leaf red nodes have a size of 8. The stopping criterion is relative residual less than 10^{-10} .

tree depth		IFMM			GC-constant		
$n \setminus \epsilon$		0.1	0.2	0.3	0.1	0.2	0.3
$(2^5)^2$	7	5	6	7	5	7	7
$(2^6)^2$	9	7	8	9	6	7	9
$(2^7)^2$	11	9	11	14	7	9	11
$(2^8)^2$	13	12	18	22	7	10	12
$(2^9)^2$	15	20	31	40	8	10	15
$(2^{10})^2$	17	33	56	74	8	12	18

In Table 4, the number of GMRES iterations is shown for different hierarchical solvers for the piecewise-constant coefficient Poisson equation. In Table 5, the results for the random coefficient Poisson equation are shown. Both of these two tables show similar performance of the IFMM and the GC-constant compared with the constant coefficient case. These results show that the GC-constant is consistently improving over the IFMM. This indicates that the GC-constant is a promising algorithm for practical engineering problems. The GC-constant leads to a number of iterations that appears bounded (in our benchmarks), implying that the preconditioned system has a bounded condition number and that the eigenvalues have a favorable distribution for GMRES.

8 Conclusion

In this paper, we modified the algorithm of the hierarchical solver in [37] and conducted a convergence analysis. Based on the convergence analysis, we proposed an improved hierarchical solver that increases the accuracy on the low-frequency components of the error. Given certain eigenvectors or an approximation thereof, the algorithm is able to exactly preserve these vectors. In Appendix A, we also propose various methods to approximately preserve these vectors, with a given accuracy. Different approaches are suggested.

As a practical choice, we considered the use of piecewise constant vectors as local approximations of the smooth eigenvectors associated with small eigenvalues. This choice applies to elliptic PDEs discretizations. This is a computationally cheap option that provides a sufficiently good approximation of the eigenvectors. Under certain assumptions, this improved solver is an optimal preconditioner. In our benchmarks for example, we observed that GMRES converges with an almost constant number of iterations, regardless of the problem size.

Acknowledgments

Funding from the “Army High Performance Computing Research Center” (AHPCRC), sponsored by the U.S. Army Research Laboratory under contract No. W911NF-07-2-0027, at Stanford, supported in part this research. The financial support is gratefully acknowledged.

A Improved convergence on low-frequency eigenvectors

Based on Theorem 2, we need the accuracy on each eigenvector to be ideally proportional to the magnitude of the corresponding eigenvalue. In order to be computationally efficient, we only improve the accuracy for a few smallest eigenvectors. In the rest of this appendix, we introduce three approaches for this purpose: **first-order one-sided projection**, **first-order symmetric projection**, and **second-order symmetric projection**. The compression is mostly related to the well-separated block A_{sw} and A_{ws} . For brevity, we denote A_{ws} by B .

During one step of compression, the well-separated interactions B and B^T are approximated by low-rank matrices. In order to improve the convergence on a few low-frequency eigenvectors: $\{e_1, e_2, \dots, e_M\}$, we need the approximations \tilde{B} and \tilde{B}^T to be more accurate on $\{[e_i]_s\}$ and $\{[e_i]_w\}$, respectively. Here $[\cdot]$ denotes a segment of a vector and subscript s and w denote the current super node and its well-separated nodes, respectively. Our objective is to find \tilde{B} such that

$$\begin{aligned} \|(B - \tilde{B})[e_i]_s\| &\leq \epsilon_1 \lambda_i \|B\| \|[e_i]_s\|, \quad \forall e_i, \quad i \leq M, \\ \|(B^T - \tilde{B}^T)[e_i]_w\| &\leq \epsilon_1 \lambda_i \|B\| \|[e_i]_w\|, \quad \forall e_i, \quad i \leq M, \end{aligned} \quad (21)$$

$$\|B - \tilde{B}\| \leq \epsilon_2 \|B\|. \quad (22)$$

Note that $\epsilon_1 \leq \epsilon_2$. The objectives mean that \tilde{B} approximates B with an accuracy of ϵ_2 . Moreover, the approximation has an improved accuracy of ϵ_1 on $\text{span}\{e_i\}_{i \leq N}$, and the accuracy is proportional to λ_i on the direction of e_i , $\forall 1 \leq i \leq N$. Let

$$X_s = \left[\frac{[e_1]_s}{\lambda_1}, \frac{[e_2]_s}{\lambda_2}, \dots, \frac{[e_M]_s}{\lambda_M} \right],$$

and

$$X_w = \left[\frac{[e_1]_w}{\lambda_1}, \frac{[e_2]_w}{\lambda_2}, \dots, \frac{[e_M]_w}{\lambda_M} \right].$$

First-order one-sided projection

We first perform a truncated SVD for $[X_s, B^T X_w]$ and obtain U_1 , the purpose of which is to approximately preserve $[X_s, B^T X_w]$. Then we keep the component of $[X_s, B^T]$ on the “direction” of U_1 and compress the rest of it via another truncated SVD.

1. Perform a truncated SVD:

$$[X_s, B^T X_w] = U_1 \Sigma_1 V_1^T + \hat{U}_1 \hat{\Sigma}_1 \hat{V}_1^T,$$

where $\|\hat{U}_1 \hat{\Sigma}_1 \hat{V}_1^T\| \leq \epsilon_1 \|[X_s, B^T X_w]\|$. Then

$$\|(I - U_1 U_1^T)[X_s, B^T X_w]\| \leq \epsilon_1 \|[X_s, B^T X_w]\|.$$

Obtain an intermediate approximation

$$\hat{B}^T = U_1 U_1^T B^T.$$

2. Perform another truncated SVD:

$$(I - U_1 U_1^T) B^T = U_2 \Sigma_2 V_2^T + E_2, \quad \text{with } \|E_2\| \leq \epsilon_2 \|B^T\|$$

Obtain the final approximation

$$\tilde{B}^T := \hat{B}^T + U_2 U_2^T (B^T - \hat{B}^T).$$

Note that for any given vector v_s and v_w , we always have

$$\begin{aligned}\|(B - \tilde{B})v_s\| &= \|B(I - U_1U_1^T)(I - U_2U_2^T)v_s\| \leq \epsilon_2\|B\|\|v_s\|, \\ \|(B^T - \tilde{B}^T)v_w\| &= \|(I - U_2U_2^T)(I - U_1U_1^T)B^Tv_w\| \leq \epsilon_2\|B^T\|\|v_w\|.\end{aligned}$$

Therefore,

$$\begin{aligned}\|(B - \tilde{B})X_s\alpha_i\| &= \|B(I - U_1U_1^T)(I - U_2U_2^T)(I - U_1U_1^T)X_s\alpha_i\| \\ &\leq \|B(I - U_1U_1^T)(I - U_2U_2^T)\|\|(I - U_1U_1^T)X_s\|\|\alpha_i\| \\ &\leq \lambda_i\epsilon_1\epsilon_2\|B\|\|[X_s, B^TX_w]\|\|\alpha_i\|.\end{aligned}$$

For $B^T - \tilde{B}^T$, we have the following estimates

$$\begin{aligned}\|(B^T - \tilde{B}^T)X_w\| &= \|(I - U_2U_2^T)(I - U_1U_1^T)B^TX_w\| \leq \epsilon_1\|[X_s, B^TX_w]\| \\ \|(B^T - \tilde{B}^T)X_w\| &= \|(I - U_2U_2^T)(I - U_1U_1^T)B^TX_w\| \leq \epsilon_2\|B^T\|\|X_w\|.\end{aligned}$$

Therefore,

$$\|(B^T - \tilde{B}^T)X_w\| \leq \min\{\epsilon_1\|[X_s, B^TX_w]\|, \epsilon_2\|B^T\|\|X_w\|\}.$$

First-order symmetric projection

In this subsection we introduce a symmetric approach to approximate preserve vectors. The approximation \tilde{B}^T is obtained via the following steps.

1. We first perform two truncated SVDs:

$$\begin{aligned}[X_s, B^TX_w] &= U_1^s\Sigma_1^s(V_1^s)^T + \hat{U}_1^s\hat{\Sigma}_1^s(\hat{V}_1^s)^T, \\ [X_w, BX_s] &= U_1^w\Sigma_1^w(V_1^w)^T + \hat{U}_1^w\hat{\Sigma}_1^w(\hat{V}_1^w)^T,\end{aligned}$$

which satisfy

$$\|\hat{U}_1^s\hat{\Sigma}_1^s(\hat{V}_1^s)^T\| \leq \epsilon_1\|[X_s, B^TX_w]\|, \quad \|\hat{U}_1^w\hat{\Sigma}_1^w(\hat{V}_1^w)^T\| \leq \epsilon_1\|[X_w, BX_s]\|.$$

Note that equivalently we have

$$\begin{aligned}\|(I - U_1^s(U_1^s)^T)[X_s, B^TX_w]\| &\leq \epsilon_1\|[X_s, B^TX_w]\|, \\ \|(I - U_1^w(U_1^w)^T)[X_w, BX_s]\| &\leq \epsilon_1\|[X_w, BX_s]\|.\end{aligned}\tag{23}$$

Then an intermediate approximation is

$$\hat{B}^T = U_1^s(U_1^s)^TB^TU_1^w(U_1^w)^T$$

2. Perform another SVD:

$$B^T - \hat{B}^T = U_2\Sigma_2V_2^T + \hat{U}_2\hat{\Sigma}_2\hat{V}_2^T,$$

where $\|\hat{U}_2\hat{\Sigma}_2\hat{V}_2^T\| \leq \epsilon_2\|B^T\|$, or, equivalently,

$$\|(I - U_2U_2^T)(B^T - \hat{B}^T)\| \leq \epsilon_2\|B^T\|.$$

Then the final approximation is obtained as follows:

$$\tilde{B}^T = \hat{B}^T + U_2U_2^T(B^T - \hat{B}^T).$$

Now we verify that the approximation \tilde{B}^T satisfies the objectives (21), and (22).

First we check whether \hat{B}^T satisfy (21).

$$\begin{aligned}(B^T - \hat{B}^T)X_w &= (B^T - U_1^s(U_1^s)^T B^T U_1^w(U_1^w)^T)X_w \\ &= (I - U_1^s(U_1^s)^T)B^T X_w + U_1^s(U_1^s)^T B^T (I - U_1^w(U_1^w)^T)X_w.\end{aligned}$$

Therefore,

$$\|(B^T - \hat{B}^T)X_w\| \leq \epsilon_1(\|[X_s, B^T X_w]\| + \|B^T\| \|[X_w, BX_s]\|).$$

We can get a similar estimate for X_s :

$$\|(B - \hat{B})X_s\| \leq \epsilon_1(\|B^T\| \|[X_s, B^T X_w]\| + \|[X_w, BX_s]\|).$$

Next, we check \tilde{B}^T . Note that (22) can be easily verified:

$$\|B^T - \tilde{B}^T\| = \|(I - U_2 U_2^T)(B^T - \hat{B}^T)\| \leq \epsilon_2 \|B^T\|$$

Now we check (21) for \tilde{B}^T :

$$\|(B^T - \tilde{B}^T)X_w\| = \|(I - U_2 U_2^T)(B^T - \hat{B}^T)X_w\| \leq \epsilon_2 \|B^T\| \|X_w\|.$$

We can also estimate it using ϵ_1 :

$$\begin{aligned}\|(B^T - \tilde{B}^T)X_w\| &= \|(I - U_2 U_2^T)(B^T - \hat{B}^T)X_w\| \\ &\leq \|(B^T - \hat{B}^T)X_w\| \\ &\leq \epsilon_1(\|[X_s, B^T X_w]\| + \|B^T\| \|[X_w, BX_s]\|)\end{aligned}$$

Therefore, we can only conclude that the upper bound for $\|(B^T - \tilde{B}^T)X_w\|$ is a minimal of

$$\|(B^T - \tilde{B}^T)X_w\| \leq \min\{\epsilon_1(\|[X_s, B^T X_w]\| + \|B^T\| \|[X_w, BX_s]\|), \epsilon_2 \|B^T\| \|X_w\|\}.$$

Similarly,

$$\|(B - \tilde{B})X_s\| \leq \min\{\epsilon_1(\|B^T\| \|[X_s, B^T X_w]\| + \|[X_w, BX_s]\|), \epsilon_2 \|B^T\| \|X_s\|\}.$$

Second-order symmetric projection

There is another symmetric approach to achieve (21) and (22). We perform two truncated SVDs:

$$X_s = U_1^s \Sigma_1^s (V_1^s)^T + \hat{U}_1^s \hat{\Sigma}_1^s (\hat{V}_1^s)^T, \quad X_w = U_1^w \Sigma_1^w (V_1^w)^T + \hat{U}_1^w \hat{\Sigma}_1^w (\hat{V}_1^w)^T,$$

which satisfy

$$\|\hat{U}_1^s \hat{\Sigma}_1^s (\hat{V}_1^s)^T\| \leq \epsilon_1 \|X_s\|, \quad \|\hat{U}_1^w \hat{\Sigma}_1^w (\hat{V}_1^w)^T\| \leq \epsilon_1 \|X_w\|.$$

Note that equivalently we have

$$\|(I - U_1^s (U_1^s)^T)X_s\| \leq \epsilon_1 \|X_s\|, \quad \|(I - U_1^w (U_1^w)^T)X_w\| \leq \epsilon_1 \|X_w\|. \quad (24)$$

Then an intermediate approximation is

$$\hat{B}^T = B^T - (I - U_1^s (U_1^s)^T)B^T(I - U_1^w (U_1^w)^T)$$

Perform another SVD:

$$B^T - \hat{B}^T = U_2 \Sigma_2 V_2^T + E_2,$$

where $\|E_2\| \leq \epsilon_2 \|B^T\|$.

Then the final approximation is obtained as follows:

$$\tilde{B}^T = \hat{B}^T + U_2 U_2^T (B^T - \hat{B}^T).$$

Now we verify that the approximation \tilde{B}^T satisfies the objectives (21), and (22).
Note that

$$\begin{aligned} B^T - \tilde{B}^T &= (I - U_2 U_2^T)(B^T - \hat{B}^T) \\ &= (I - U_2 U_2^T)(I - U_1^s (U_1^s)^T) B^T (I - U_1^w (U_1^w)^T). \end{aligned}$$

We start with (22):

$$\|B^T - \tilde{B}^T\| = \|E_2\| \leq \epsilon_2 \|B^T\|. \quad (25)$$

Then we verify (21). For the “s” component, we have for $v_s = X_s \alpha$

$$\begin{aligned} \|(B - \tilde{B})X_s\| &= \|(I - U_1^w (U_1^w)^T) B (I - U_1^s (U_1^s)^T) (I - U_2 U_2^T) X_s\| \\ &= \|(I - U_1^w (U_1^w)^T) B (I - U_1^s (U_1^s)^T) (I - U_2 U_2^T) (I - U_1^s (U_1^s)^T) X_s\| \\ &\leq \|B - \tilde{B}\| \|(I - U_1^s (U_1^s)^T) X_s\| \\ &\leq \epsilon_1 \epsilon_2 \|B^T\| \|X_s\|. \end{aligned}$$

Note that we should always have

$$(I - U_1^s (U_1^s)^T) (I - U_2 U_2^T) = (I - U_1^s (U_1^s)^T) (I - U_2 U_2^T) (I - U_1^s (U_1^s)^T)$$

no matter whether U_2 and U_1^s are orthogonal or not.

In particular, for one of the low-frequency eigenvectors $[e_i]_s$, $1 \leq i \leq N$, we have

$$[e_i]_s = X_s \alpha_i, \quad \text{where } \alpha_i = \begin{pmatrix} 0 \\ \vdots \\ \lambda_i \\ \vdots \\ 0 \end{pmatrix}.$$

Then

$$\|(B - \tilde{B})[e_i]_s\| \leq \lambda_i \epsilon_1 \epsilon_2 \|B^T\| \|X_s\|, \quad 1 \leq i \leq N.$$

For “w”, it is similar. We take $v_w = X_w \alpha$

$$\begin{aligned} \|(B^T - \tilde{B}^T)X_w\| &= \|(I - U_2 U_2^T) (I - U_1^s (U_1^s)^T) B^T (I - U_1^w (U_1^w)^T) X_w\| \\ &= \|(I - U_2 U_2^T) (I - U_1^s (U_1^s)^T) B^T (I - U_1^w (U_1^w)^T) (I - U_1^w (U_1^w)^T) X_w\| \\ &\leq \|B^T - \tilde{B}^T\| \|(I - U_1^w (U_1^w)^T) X_w\| \\ &\leq \epsilon_1 \epsilon_2 \|B^T\| \|X_w\|. \end{aligned}$$

Then

$$\|(B^T - \tilde{B}^T)[e_i]_w\| \leq \lambda_i \epsilon_1 \epsilon_2 \|B^T\| \|X_w\|, \quad 1 \leq i \leq N.$$

These approaches approximately preserve vectors for hierarchical solvers. The efficiency of them is to be explored in future work.

References

- [1] S. AMBIKASARAN AND E. DARVE, *An $\mathcal{O}(n \log n)$ fast direct solver for partial hierarchically semi-separable matrices*, Journal of Scientific Computing, 57 (2013), pp. 477–501.
- [2] S. AMBIKASARAN AND E. DARVE, *The inverse fast multipole method*, arXiv preprint arXiv:1407.1572, (2014).
- [3] A. AMINFAR, S. AMBIKASARAN, AND E. DARVE, *A fast block low-rank dense solver with applications to finite-element matrices*, Journal of Computational Physics, 304 (2016), pp. 170–188.
- [4] A. AMINFAR AND E. DARVE, *A fast, memory efficient and robust sparse preconditioner based on a multifrontal approach with applications to finite-element matrices*, International Journal for Numerical Methods in Engineering, (2016).
- [5] E. ANDERSON, Z. BAI, C. BISCHOF, S. BLACKFORD, J. DEMMEL, J. DONGARRA, J. DU CROZ, A. GREENBAUM, S. HAMMERLING, A. MCKENNEY, ET AL., *LAPACK Users' guide*, vol. 9, SIAM, 1999.
- [6] M. BEBENDORF, *Why finite element discretizations can be factored by triangular hierarchical matrices*, SIAM Journal on Numerical Analysis, 45 (2007), pp. 1472–1494.
- [7] M. BEBENDORF, *Hierarchical matrices*, Springer, 2008.
- [8] M. BEBENDORF, M. BOLLHÖFER, AND M. BRATSCH, *Hierarchical matrix approximation with blockwise constraints*, BIT Numerical Mathematics, 53 (2013), pp. 311–339.
- [9] M. BEBENDORF, M. BOLLHÖFER, AND M. BRATSCH, *On the spectral equivalence of hierarchical matrix preconditioners for elliptic problems*, Mathematics of Computation, 85 (2016), pp. 2839–2861.
- [10] S. BÖRM, L. GRASEDYCK, AND W. HACKBUSCH, *Introduction to hierarchical matrices with applications*, Engineering Analysis with Boundary Elements, 27 (2003), pp. 405–422.
- [11] J. H. BRAMBLE, *Multigrid methods*, vol. 294, CRC Press, 1993.
- [12] A. BRANDT, *Algebraic multigrid theory: The symmetric case*, Applied Mathematics and Computation, 19 (1986), pp. 23–56.
- [13] A. BRANDT, S. MCCORUICK, AND J. HUGE, *Algebraic multigrid (AMG) for sparse matrix equations*, Sparsity and its Applications, (1985), p. 257.
- [14] W. L. BRIGGS, V. E. HENSON, AND S. F. MCCORMICK, *A multigrid tutorial*, SIAM, 2000.
- [15] S. CHANDRASEKARAN, P. DEWILDE, M. GU, W. LYONS, AND T. PALS, *A fast solver for HSS representations via sparse matrices*, SIAM Journal on Matrix Analysis and Applications, 29 (2006), pp. 67–81.
- [16] S. CHANDRASEKARAN, M. GU, AND T. PALS, *A fast ULV decomposition solver for hierarchically semiseparable representations*, SIAM Journal on Matrix Analysis and Applications, 28 (2006), pp. 603–622.
- [17] E. CORONA, P.-G. MARTINSSON, AND D. ZORIN, *An direct solver for integral equations on the plane*, Applied and Computational Harmonic Analysis, 38 (2015), pp. 284 – 317, doi:<http://dx.doi.org/10.1016/j.acha.2014.04.002>, <http://www.sciencedirect.com/science/article/pii/S1063520314000529>.
- [18] P. COULIER, H. POURANSARI, AND E. DARVE, *The inverse fast multipole method: using a fast approximate direct solver as a preconditioner for dense linear systems*, arXiv preprint arXiv:1508.01835, (2015).

- [19] E. DARVE, *The fast multipole method: numerical implementation*, Journal of Computational Physics, 160 (2000), pp. 195–240.
- [20] W. FONG AND E. DARVE, *The black-box fast multipole method*, Journal of Computational Physics, 228 (2009), pp. 8712–8725.
- [21] A. GILLMAN, P. M. YOUNG, AND P.-G. MARTINSSON, *A direct solver with $o(n)$ complexity for integral equations on one-dimensional domains*, Frontiers of Mathematics in China, 7 (2012), pp. 217–247.
- [22] L. GREENGARD, D. GUEYFFIER, P.-G. MARTINSSON, AND V. ROKHLIN, *Fast direct solvers for integral equations in complex three-dimensional domains*, Acta Numerica, 18 (2009), pp. 243–275.
- [23] L. GREENGARD AND V. ROKHLIN, *A fast algorithm for particle simulations*, Journal of Computational Physics, 73 (1987), pp. 325–348.
- [24] G. GUENNEBAUD, B. JACOB, ET AL., *Eigen v3*. <http://eigen.tuxfamily.org>, 2010.
- [25] W. HACKBUSCH, *A sparse matrix arithmetic based on \mathcal{H} -matrices. part i: Introduction to \mathcal{H} -matrices*, Computing, 62 (1999), pp. 89–108.
- [26] W. HACKBUSCH, *Multi-grid methods and applications*, vol. 4, Springer Science & Business Media, 2013.
- [27] W. HACKBUSCH AND S. BÖRM, *\mathcal{H}^2 -matrix approximation of integral operators by interpolation*, Applied Numerical Mathematics, 43 (2002), pp. 129–143.
- [28] K. L. HO AND L. YING, *Hierarchical interpolative factorization for elliptic operators: differential equations*, Communications on Pure and Applied Mathematics, (2015).
- [29] K. L. HO AND L. YING, *Hierarchical interpolative factorization for elliptic operators: integral equations*, Communications on Pure and Applied Mathematics, (2015).
- [30] W. Y. KONG, J. BREMER, AND V. ROKHLIN, *An adaptive fast direct solver for boundary integral equations in two dimensions*, Applied and Computational Harmonic Analysis, 31 (2011), pp. 346–369.
- [31] R. LI AND Y. SAAD, *Divide and conquer low-rank preconditioners for symmetric matrices*, SIAM Journal on Scientific Computing, 35 (2013), pp. A2069–A2095.
- [32] A. NAPOV AND X. S. LI, *An algebraic multifrontal preconditioner that exploits the low-rank property*, Numerical Linear Algebra with Applications, 23 (2016), pp. 61–82.
- [33] A. NAPOV AND Y. NOTAY, *An algebraic multigrid method with guaranteed convergence rate*, SIAM Journal on Scientific Computing, 34 (2012), pp. A1079–A1109.
- [34] N. NISHIMURA, *Fast multipole accelerated boundary integral equation methods*, Applied Mechanics Reviews, 55 (2002), pp. 299–324.
- [35] Y. NOTAY, *An aggregation-based algebraic multigrid method*, Electronic Transactions on Numerical Analysis, 37 (2010), pp. 123–146.
- [36] F. PELLEGRINI AND J. ROMAN, *SCOTCH: A software package for static mapping by dual recursive bipartitioning of process and architecture graphs*, in High-Performance Computing and Networking, Springer, 1996, pp. 493–498.
- [37] H. POURANSARI, P. COULIER, AND E. DARVE, *Fast hierarchical solvers for sparse matrices*, arXiv preprint arXiv:1510.07363, (2015).
- [38] H. POURANSARI AND E. DARVE, *Optimizing the adaptive fast multipole method for fractal sets*, SIAM Journal on Scientific Computing, 37 (2015), pp. A1040–A1066.

- [39] J. RUGE AND K. STÜBEN, *Algebraic multigrid*, vol. 3, SIAM, Philadelphia, 1987.
- [40] Y. SAAD, *Iterative methods for sparse linear systems*, SIAM, 2003.
- [41] K. STÜBEN, *A review of algebraic multigrid*, Journal of Computational and Applied Mathematics, 128 (2001), pp. 281–309.
- [42] D. A. SUSHNIKOVA AND I. V. OSELEDETS, “*Compress and eliminate*” solver for symmetric positive definite sparse matrices, arXiv preprint arXiv:1603.09133, (2016).
- [43] U. TROTTEBERG, C. W. OOSTERLEE, AND A. SCHULLER, *Multigrid*, Academic Press, 2000.
- [44] J. XIA, S. CHANDRASEKARAN, M. GU, AND X. S. LI, *Fast algorithms for hierarchically semiseparable matrices*, Numerical Linear Algebra with Applications, 17 (2010), pp. 953–976.
- [45] J. XU, *Iterative methods by space decomposition and subspace correction*, SIAM Review, 34 (1992), pp. 581–613.
- [46] L. YING, G. BIROS, AND D. ZORIN, *A kernel-independent adaptive fast multipole algorithm in two and three dimensions*, Journal of Computational Physics, 196 (2004), pp. 591–626.